
The NIST Risk Management Framework: Problems and recommendations

Received (in revised form): 14th August, 2017



Don Maclean

is Chief Cyber Security Technologist for DLT and formulates and executes cyber security portfolio strategy, speaks and writes on security topics, and socialises his company's cyber security portfolio. He has nearly 30 years' experience working with US federal agencies. Before joining DLT in 2015, he managed security programmes for numerous US federal clients, including DOJ, DOL, FAA, FBI and the Treasury Department. This experience allowed him to work closely with the NIST Risk Management Framework featured in this paper, and to understand its strengths and weaknesses. In addition to CISSP, PMP, CEH and CCSK certificates, he holds a BA in music from Oberlin College and Conservatory, an MS in information security from Brandeis Rabb School, and is nearing completion of his second Bachelor's in mathematics. An avid musician, he organises a concert for charity every year, and has been known to compete in chess and Shogi (Japanese chess) tournaments, both in person and online.

Chief Cybersecurity Technologist, DLT, 2411 Dulles Corner Park, Herndon, VA 20171, USA
Tel: 571-346-1854; E-mail: don.maclean@dlit.com

Abstract Cyber security assessment initiatives and frameworks abound in the US government, but their effectiveness is inconsistent. The most important law from which these frameworks and assessments arose is the Federal Information Systems Management Act (FISMA), passed in 2002, and updated as the Federal Information Systems Modernization Act in 2014. The law's broad scope included a mandate to the US National Institute of Standards and Technology (NIST), charging it to create methods and standards to assess and optimise the cyber security posture of US government agencies. NIST's flagship methodology, Risk Management Framework (RMF), is comprehensive and fundamentally sound, but years of experience have exposed many flaws — some stemming from lack of proper adoption and execution, some from unintended consequences, and others arising from the relentless pace of innovation in technology. This paper examines the RMF's weaknesses, and offers recommendations for improvement.

KEYWORDS: National Institute of Science and Technology, Risk Management Framework, Plan of Action and Milestones, Federal Information Systems Management Act, risk assessment, Authority to Operate

BACKGROUND

The shortcomings of the Risk Management Framework (RMF) of the US National Institute of Standards and Technology (NIST) reared their ugly head in the 2015 breach of the US Office of Personnel Management (OPM). In that intrusion,

personnel records of 21.5m government staff and contractors — many of whom had high-level clearances — were purloined from an agency that had received satisfactory marks for cyber security.

Many components of the US government — including NIST itself — have set forth

other initiatives and frameworks to mitigate the evident weakness of NIST's RMF methodology. These include:

- The intentionally short-lived cyber security 'Sprint' of 2015;¹
- Cybersecurity National Action Plan (CNAP);²
- Continuous Diagnostics and Mitigation (CDM) programme;³
- NIST's Cyber Security Framework (CSF) (not to be confused with the NIST RMF);⁴
- Other initiatives and presidential directives.

Although obscured by a maze of bureaucracy and an olio of acronyms, the need for better cyber security drove each effort. Space does not permit examination of each of these programmes; however, the paper will describe the problems with the Federal Information Systems Management Act (FISMA) and the NIST RMF, and provide recommendations for their improvement. The recommendations are intended for those seeking to adopt the NIST RMF, so they can benefit from its virtues, but avoid its vices.

The problem descriptions and the recommendations are offered in the spirit of constructive criticism. NIST's work is consistently thorough and conscientious, and US government security programmes, although still wanting, would be even more deficient without it. Finally, please note that at the time of writing, initiatives arising from the Trump Administration are too embryonic for reasonable assessment.

Initiatives and frameworks spring from good intentions, but major overhauls are necessary. The US government must recognise and accept the weaknesses that have led to major problems, and take steps to rectify them. A significant principle is to accept that cyber security is fundamentally different from other technologies: cyber security tools are at their best when new, and quantifying success in cyber security is

difficult. Therefore, it is essential to provide the means for rapid early adoption, and to work toward quantifiable assessment of security programme efficacy.

Achieving these goals, however, requires acknowledgment and rectification of the problems plaguing security programmes in the federal government. These include:

- Conflicts of interest;
- Plan of Action and Milestone (POA&M) abuse;
- Lack of incentives;
- Excessive emphasis on compliance compounded by burdensome documentation requirements;
- Risk scoring: quantitative vs qualitative scoring;
- Categorisation: high-water mark;
- Impact assessment: focus on victim, not organisation.

FISMA

In the US, the Federal Information Systems Management Act of 2002 (FISMA)⁵ is the 'mother of all cybersecurity laws'. Fifteen years from its inception, FISMA's efficacy is debatable. Without it, cyber security would likely receive even less funding attention than it does now, and the status of cyber security programmes might well be more woeful; however, FISMA has also diverted resources away from tangible measures and toward compliance efforts such as documentation, procurement practices, budget allocation and the like. FISMA's broad scope is admirable — it covers all major areas of concern — but assessing and demonstrating compliance is excessively burdensome. For example, in one federal agency, nearly 70 per cent of the cyber security budget goes to compliance paperwork, to the detriment of more concrete cyber security measures. Agencies can easily spend more time and effort documenting a problem than fixing it. The process often looks like this:

- Discover a security weakness;
- Document the weakness and how it was found;
- Assess the impact of the weakness;
- Determine if the weakness has been exploited;
- Identify remediation methods and possible compensating controls;
- Estimate the cost of remediation;
- Evaluate alternative methods of mitigation or compensation;
- Assign responsibility for the problem and its repair;
- Test remediation procedures in a lab;
- Seek approval for the change through a Change Control Board or equivalent (many of which meet only monthly);
- Write detailed technical procedures to implement the measure;
- Plan back-out procedures if the remediation effort go awry;
- Announce the change to users;
- Respond to objections from users or other staff;
- Fix the problem.

This writer personally observed one incident in which the ancillary work consumed nearly 40 hours, while the process of reparation took less than five minutes. By comparison, the attackers' checklist looks more this:

- Discover a security weakness;
- Exploit the security weakness to steal data;
- Sell the data;
- Vacation in Tahiti.

Compliance, bureaucracy and documentation have a necessary and worthwhile place in cyber security, but they have become disproportionately prominent priorities, to the detriment of the very security posture they seek to ensure. Hackers move quickly and carry no bureaucratic overhead, giving them a large and growing lead in the race. To narrow the gap, and ultimately put the adversary in the rear-view mirror,

agencies must be able to move more quickly. Achieving this agility will require a tectonic shift in the culture of the massive organisation called the US government, but continuance of current practices will ensure defeat.

Paradoxically, the overemphasis on compliance has engendered few effective incentives to achieve or maintain a FISMA-compliant security posture. Enforcement measures are feckless — the infamous OPM breach resulted in precisely one dismissal — and positive incentives are anaemic.

The law broadly requires a framework for information security management and mandates NIST to develop specific standards, guidelines and technical requirements. Agency officials must conduct an annual review of their organisation's security posture and report the result to the Office of Management and Budget (OMB), which compiles and publishes the results.⁶

Nonetheless, FISMA delegates most of the 'heavy lifting' to NIST, whose framework⁷ includes three major components:

1. **Information Security Automation Program (ISAP).** ISAP aims to automate validation and remediation of security deficiencies and to minimise human involvement in compliance activities. Supporting objectives include standardisation of vulnerability data (at rest and in transit), baselines for IT products, systems assessment and reporting, and the use of metrics to assess risk and vulnerability.
2. **National Vulnerability Database (NVD).** Using the Security Content Automation Protocol (SCAP) to track and manage vulnerabilities, the NVD categorises vulnerabilities quantitatively. Vulnerability databases are available in numerous formats, such as JSON and the Common Vulnerability Enumeration (CVE) format.

FISMA implementation

This is the heart and soul of FISMA, around which most of the NIST standards documents revolve. The large-scale architecture of FISMA implementation is the **NIST Risk Management Framework (RMF)** consisting of:

- **Inventory** of systems, primarily hardware and software.
- **System categorisation**, which determines the selection and stringency of security measures.
- **Security controls**, strictly defined by extensive NIST documentation.
- **Risk assessment**, which summarises the risk facing each system and underlies budget allocation decisions.
- **System security plan** — an ill-named document that details security controls already in place, rather than those planned for implementation.
- **Certification and accreditation/assessment and authorisation** — the formal process of assessing and documenting the security posture of each system, with the goal of attaining an Authority to Operate (ATO) for every system.
- **Plan of action and milestones** — the document that details the weaknesses of each system and the plans for mitigating them.
- The paper will also examine **conflicts of interest**, an unfortunate but unintended consequence of FISMA implementation.

1. **Inventory.** One might assume that agencies would track such fundamental information as a matter of course. FISMA, however, was promulgated in 2002, and in the OMB report of 2016 hardware and software asset management still emerges as a major security deficiency for many agencies. Obviously, it is very difficult to secure assets whose very existence is unknown;

2. **System categorisation.** The system categorisation drives many security decisions, including budget allocation. The NIST framework requires assignment of an impact level — low, moderate or high — to the confidentiality, integrity and availability of each type of data under the agencies' purview.^{8,9} NIST provides several ancillary documents to ensure correct execution of the process, such as the snappily titled cliffhanger 'SP800-60 Guide for Mapping Types of Information and Information Systems to Security Categories, Vol. 1'¹⁰ and the heart-pounding sequel 'Appendices to Guide for Mapping Types of Information and Information Systems to Security Categories, Vol. 2'.¹¹

A key aspect of — and important problem with — system categorisation is the 'high-water mark' approach: if a system's confidentiality is assessed at 'high', then every other aspect of the system in question is treated as 'high' as well. For example, some data may be very confidential, but used infrequently. In the high-water mark approach, however, the agency is required to implement controls for availability and integrity at the 'high' level. Data used once a week must be made available without fail, 24 hours per day, seven days per week, 365 days per year.

The intent of the high-water mark approach is to ensure maximum security for each system; however, the effect can be quite different. Agencies concerned about costs recognise that a 'high' designation will incur expenses for unnecessary security measures, such as high availability. Agencies might, then, undercategorise a system to avoid budget issues. Alternatively, they will categorise the system correctly, but write a series of waivers to nullify needless expenses. This approach, when overused, can become a bad habit: waivers can replace security all too easily.

There are two other flaws in the security categorisation process: misleading ‘quantification’ (also a problem in risk assessment), and confusion over whom the impact really affects. These issues also arise in the risk assessment process, discussed below.

To assess the impact of data loss, agencies are directed to rate the impact of a data loss in terms of security’s ‘golden triangle’ — confidentiality, integrity and availability — using a rating scale that typically goes from 1 to 5. Ratings are a matter of professional judgment, supposition and educated guesswork, but the numeric scale creates the (false) appearance of a strict, quantitative evaluation. This creates ample, or excessive, room for error, misinterpretation or manipulation of the impact assessment.

Moreover, the impact assessment almost universally addresses the impact of data loss on the *organisation*, but not on the *victim* of a breach. Refer again to the OPM hack: the true victims were the individuals whose data is now being bought and sold for profit on the Dark Web, but the impact assessment almost certainly addressed the effect on OPM, not on those whose records were purloined;

3. Security controls:

Control descriptions. NIST’s compendium of security controls, ‘SP800-53r4, Security and Privacy Controls for Federal Information Systems and Organizations’¹² describes hundreds of aspects of computer system security in exquisite detail. Each control belongs to a larger control family, and most include control ‘enhancements’, which may be either required or optional. In addition, controls often offer a measure of latitude to the organisation implementing them.

Figure 1 shows a sample control from SP800-53r4. The sample belongs to the ‘Audit and Accountability’

family, abbreviated ‘AU’. This control includes two *enhancements*, designated, with startling creativity, as paragraphs (1) and (2). The first is required for high-impact systems; the second is optional for any system. Specifics of the control requirements — in this case, the granularity of the time-stamp records — are up to the agency.

Control assessments. For each control, there is a corresponding set of assessment procedures, spelled out in SP800-53A;¹³ an excerpt appears in Figure 2. Note that for each control, one or more assessment method may be applicable: Examine, Interview, or Test. The ‘Examine’ method refers to the review of documentation on a control, the ‘Interview’ method involves interviewing appropriate staff to evaluate the control’s implementation, and the ‘Test’ method requires testing a control to ensure that it is implemented correctly and functioning as intended.

Notice that two of three methods — ‘Examine’ and ‘Interview’ — are *indirect* evaluations of a control’s effectiveness. Only the ‘Test’ method requires direct observation of the control and its efficacy.

In this writer’s observation, examining documentation and interviewing personnel are much easier to accomplish than testing a control, yet they provide ample material for bulked-up documentation ostensibly demonstrating deep scrutiny of a control.

The emphasis on ‘Interview’ and ‘Examine’ requirements contributes to cumbersome documentation resulting from a security assessment;

4. Risk assessment

NIST’s guidance¹⁴ here is fundamentally sound, but the use of numeric scoring methods is problematic. The basic steps in a risk assessment are:

- a. Identify threat sources and events.
- b. Identify vulnerabilities and predispositions.

AU-8 TIME STAMPS

Control: The information system:

- a. Uses internal system clocks to generate time stamps for audit records; and
- b. Records time stamps for audit records that can be mapped to Coordinated Universal Time (UTC) or Greenwich Mean Time (GMT) and meets [*Assignment: organization-defined granularity of time measurement*].

Supplemental Guidance: Time stamps generated by the information system include date and time. Time is commonly expressed in Coordinated Universal Time (UTC), a modern continuation of Greenwich Mean Time (GMT), or local time with an offset from UTC. Granularity of time measurements refers to the degree of synchronization between information system clocks and reference clocks, for example, clocks synchronizing within hundreds of milliseconds or within tens of milliseconds. Organizations may define different time granularities for different system components. Time service can also be critical to other security capabilities such as access control and identification and authentication, depending on the nature of the mechanisms used to support those capabilities. Related controls: AU-3, AU-12.

Control Enhancements:

(1) TIME STAMPS | SYNCHRONIZATION WITH AUTHORITATIVE TIME SOURCE

The information system:

- (a) Compares the internal information system clocks [*Assignment: organization-defined frequency*] with [*Assignment: organization-defined authoritative time source*]; and
- (b) Synchronizes the internal system clocks to the authoritative time source when the time difference is greater than [*Assignment: organization-defined time period*].

Supplemental Guidance: This control enhancement provides uniformity of time stamps for information systems with multiple system clocks and systems connected over a network.

(2) TIME STAMPS | SECONDARY AUTHORITATIVE TIME SOURCE

The information system identifies a secondary authoritative time source that is located in a different geographic region than the primary authoritative time source.

References: None.

Priority and Baseline Allocation:

P1	LOW AU-8	MOD AU-8 (1)	HIGH AU-8 (1)
----	----------	--------------	---------------

Figure 1: Sample control from SP800-53r4, Security and Privacy Controls for Federal Information Systems and Organizations

POTENTIAL ASSESSMENT METHODS AND OBJECTS:

Examine: [*SELECT FROM:* Audit and accountability policy; procedures addressing time stamp generation; information system design documentation; information system configuration settings and associated documentation; information system audit records; other relevant documents or records].

Interview: [*SELECT FROM:* Organizational personnel with information security responsibilities; system/network administrators; system developers].

Test: [*SELECT FROM:* Automated mechanisms implementing time stamp generation].

Figure 2: Excerpt from SP800-53A control assessment procedures

- c. Determine the likelihood, impact and risk to the organisation.
- d. Communicate and share information from the assessment.
- e. Monitor the risk factors and update the assessment periodically.

NIST guidance allows a government agency to choose a quantitative, qualitative or semi-quantitative method of scoring vulnerabilities, risks, adversary capability, etc. As with the impact analysis, using numbers to express subjective judgments can produce misleading results. We associate objectivity and certainty with numbers, but these numbers reflect judgments and perceptions reflect subjective and uncertain levels of perception. The use of a common formula — risk = likelihood x impact — serves to further obscure the underlying data;

5. **System security plan.** The word ‘plan’ for this document¹⁵ has caused much confusion. The ‘plan’ is the detailed and comprehensive report of the assessment of every security control: its implementation, the rationale for its use, its effectiveness, its deficiency. Note that the ‘plan’ is written *after* the controls are implemented and assessed, whereas it should be written *prior* to implementation and assessment;
6. **Certification and accreditation/assessment and authorisation.** The Certification and Accreditation (C&A)¹⁶ process, now rebranded as Assessment and Authorization (A&A), is the broad, overarching process of assessing the security posture of a system, generating a long list of hefty documents to describe that posture, and seeking authorisation from appropriate officials within a government organisation. Each agency prioritises the assessment differently, and each reviews the resulting documents with a different level of scrutiny.

In some organisations, technical and non-technical personnel read each

and every word of a huge collection of security assessment documents, and mandate appropriate action based on the findings. In others, the assessor provides the same documents, but the signatories examine them cursorily, if at all, before signing them.

It is essential for government agencies to review security assessments carefully, to recognise security problems and ensure their rapid and effective resolution. This leads to the next topic: the Plan of Action and Milestones;

7. Plan of action and milestones.

The mechanism for tracking security weaknesses is the Plan of Action and Milestones (POA&M). Misuse of the POA&M may be the single biggest, but unacknowledged, problem facing US government security programmes.

A POA&M details each security problem, describes possible methods of remediation, and establishes deadlines and cost estimates for remediation. Unfortunately, the POA&M also serves as a vehicle for budget manipulation, and can also be a tool for delaying security repairs.

Cost estimation. In a POA&M, cost estimates are often inaccurate, sometimes from lack of guidance, sometimes for more disingenuous reasons.

Since accurate estimation is difficult, the natural tendency is to exaggerate the cost of a remediation action: better to ask for money and not need it, than the reverse. These estimates are crucial to the proper execution of a security programme, but NIST provides little guidance in their creation. Consequently, many estimates are essentially arbitrary, or based on external factors such as political expediency.

Cost estimates also suffer from wilful exaggeration, which can serve an internecine political purpose or facilitate budget manipulation. For example, at one organisation, a password policy document

required an update. The requirement — to update two documents with one sentence to specify password requirements — would have taken a few minutes to write, and perhaps five minutes to undergo review and approval. The associated POA&M estimated that the fix would take three months and would cost US\$25,000. No one writing the estimate believed it: it was promulgated to seek funding for other security activities whose budgets were short. (The US\$25,000 formed part of an omnibus budget request, which hid the underlying details of the numbers.)

Schedules and deadlines. Enforcement of POA&M deadlines is also deficient. An organisation will typically create a deadline structure such as this for fixing problems: severe problems, one month; moderate problems, two months; minor issues, three months. Too often, though, the organisation will miss a deadline, either from lack of funds, lack of available resources, or simple neglect of the task. To maintain a positive image in the eyes of superiors, the organisation's management will often just change the deadlines, to show there are no overdue issues. This scenario repeats indefinitely: as the new deadline approaches, it is moved again. In one case, a problem persisted for nearly eight years before it was addressed — a process that took approximately 15 minutes to execute.

So, abuse of deadlines and inaccurate estimates — disingenuous or otherwise — constitute a major stumbling block to the proper execution of a security programme;

8. **Conflicts of interest.** Another major problem with government security programmes arises from the contractual and business relationships in which they take place. Typically, an agency will contract with a private firm to assess, oversee and maintain the security of an agency. This sounds quite sensible on the

surface, until you realise that the company in question must often produce findings and write documents that will make the agency look bad. This creates conflicting incentives: the stated goal is to discover and report security problems, but security firms endure significant contrary pressure to avoid staining the reputation of their customer.

Compounding this pressure is another typical contractual structure, in which the security programme is part of an umbrella contract to manage all aspects of an agency's information technology systems. In this very common scenario, the security group finds itself unearthing problems with the hiring agency and its *own* company.

As a consequence, too many problems go undocumented, go unrepaired, are repaired but not reported, or are simply neglected. A company hired to oversee security activities has no authority to revise erroneous or bad-faith cost estimates — which they can be pressurised to initiate — nor can they enforce POA&M deadlines. In fact, they have material incentives to do the opposite, to remain in favour with those who sign their pay cheques.

The most forceful action a security assessor can take, when assessing a government agency's security posture, is to recommend that a system be denied the Authority to Operate (ATO). In theory, this action ensures that systems with egregiously poor security do not go into operation until their security problems are repaired; it is almost never used, for the reasons described above.

RECOMMENDATIONS

Although these recommendations are tailored primarily for the US government, they should also be useful to anyone looking to adopt the NIST Risk Management Framework, or similar discipline, to manage

and assess an organisational cyber security programme.

1. **Conflict of interest**

- a. Companies performing security assessments should be contracted by an agency separate from the agency under scrutiny. GAO, DHS or OMB are possibilities, but regardless of the centralised agency, it needs to be separate from the agency being examined.
- b. Moreover, the company performing the assessment must not also have a contract to perform related IT functions — code development, systems management, cloud migration, etc. — on the systems under assessment.
- c. A related regulatory endeavour, FedRAMP, seeks to consolidate security assessments of vendors providing cloud services. Under FedRAMP, a security assessment must be carried out by an authorised company, known as a third-party assessment organisation (3PAO). The 3PAO requirement contributes two important elements to assessment integrity:
 - i. It ensures the independence of the 3PAO.
 - ii. It imposes standards of competence on the companies performing the assessment.

The 3PAO model — an independent assessor held to consistent standards of competence — would improve the integrity of the assessment process by eliminating conflicts of interest, and ameliorating the issue of assessors who lack sufficient training for the task.

2. **POA&M abuse.** The POA&M provides a vehicle for undisciplined scheduling and budget manipulation. These abuses, described above, will

continue in the absence of stronger oversight. Tangible enforcement by a centralised agency such as GAO would be ideal, but would entail a culture shift too titanic to be viable.

A reporting mechanism, such as the one in place for the DHS Continuous Diagnostics and Mitigation (CDM) programme, might be possible, however. It would reduce the incidence of deadline changes, since agencies curating POA&Ms would have no incentive to alter schedules. Making budget estimates known outside the fiefdom to which they apply would give an incentive to create cost projections more realistic than the current norm.

3. **Lack of incentives.** Government workers receive little or no extra money, time off, or other incentive for high levels of performance in protecting the systems and data in their care. Conversely, there is little disincentive for poor performance. Most agencies take security duties seriously, motivated by a sense of duty and desire to serve and protect the country. Some, however, do not.

In both cases, a bonus structure for performance, as well as more stringent accountability for dereliction of duty, would materially improve security. Unfortunately, such structures are very difficult to implement for government employees.

Contractors play a significant role in administration and security of government IT systems, however. It is relatively simple to create contracts that include bonuses for good performance and penalties for lack of proper execution. This approach is a potentially promising way to deal with the incentive issue, especially if an unbiased auditor assesses the contractor's performance.

4. **Excessive emphasis on compliance compounded by burdensome documentation requirements.** As described in the summary above,

assessing, documenting and demonstrating compliance can easily be more expensive and time-consuming than actually being compliant. This issue is widely recognised, but true solutions are hard to devise, still harder to implement. The problem is systemic and has multiple causes, and therefore requires a variegated solution. Reducing conflicts of interest and POA&M abuse, performance-based incentives in contracts, and the other reforms recommended here are all ingredients in the solution to this problem.

5. **Risk scoring: quantitative vs qualitative scoring.** ‘Quantitative’ risk scoring is misleading. Numeric estimates based on intuition and judgment measure nothing tangible, but convey a false impression of precision and scientific rigour. It is true that the NIST methodology¹⁷ provides ‘qualitative’ and ‘semi-quantitative’ scoring using verbal descriptions to supplement or replace numbers.
6. **Categorisation.** The ‘high-water mark’ approach — which mandates excessive measures in many cases — is easy to rescind; many agencies are doing so in practice. Numeric scoring, however, is embedded culturally and organisationally. This method provides a false sense of certainty, which subsumes its misleading nature. NIST would be well advised to endorse only qualitative measures of scoring, both for system categorisation and for risk assessment.
7. **Impact assessment.** Impact assessments evaluate the effect of intrusions on the organisation, primarily the organisation’s ability to perform its mission. Many government missions are matters of life and death, so it is right and proper to anticipate and plan for disruptions. Many government agencies, however, are responsible for data whose loss affects individuals more than the agency itself; OPM is a clear example. When assessing the potential impact of a breach, then, it

is essential to look at the impact from all points of view, especially those most likely to suffer negative consequences.

Conclusion

The NIST RMF is a formidable body of work that has greatly enhanced cyber security in the US federal government, and has been adopted or used in modified form in private industry and other governments. As with any work of broad scope, its consequences are difficult to foresee, especially given the complexity of an organisation as large as the US federal government, and the inherently unpredictable nature of technology in general and cyber security in particular. The suggestions offered here are intended to illuminate the issues that have emerged since the RMF’s inception, and point the way to improve an already excellent methodology for managing risk.

References

1. Cybersecurity Sprint, available at <https://obamawhitehouse.archives.gov/blog/2015/06/17/fact-sheet-enhancing-and-strengthening-federal-government-s-cybersecurity> (accessed 21st August, 2017).
2. Cybersecurity National Action Plan (CNAP), available at <https://obamawhitehouse.archives.gov/the-press-office/2016/02/09/fact-sheet-cybersecurity-national-action-plan> (accessed 21st August, 2017).
3. Continuous Diagnostics and Mitigation (CDM), available at <https://www.dhs.gov/cdm> (accessed 21st August, 2017).
4. Framework for Improving Critical Infrastructure Cybersecurity, available at <https://www.nist.gov/sites/default/files/documents/cyberframework/cybersecurity-framework-021214.pdf> (accessed 21st August, 2017).
5. Federal Information Systems Management Act (FISMA), available at <http://csrc.nist.gov/drivers/documents/FISMA-final.pdf> (accessed 21st August, 2017).
6. FISMA 2014 Annual Report to Congress, Fiscal Year 2016, available at https://www.whitehouse.gov/sites/whitehouse.gov/files/briefing-room/presidential-actions/related-omb-material/fy_2016_fisma_report%20to_congress_official_release_march_10_2017.pdf (accessed 21st August, 2017).
7. NIST Risk Management Framework, available at

- <http://csrc.nist.gov/groups/SMA/fisma/framework.html> (accessed 21st August, 2017).
8. Federal Information Processing Standards Publication 199 'Standards for Security Categorization of Federal Information and Information Systems', available at <http://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.199.pdf> (accessed 21st August, 2017).
 9. Federal Information Processing Standards Publication 200 'Minimum Security Requirements for Federal Information and Information Systems', available at <http://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.200.pdf> (accessed 21st August, 2017).
 10. Volume I: Guide for Mapping Information and Information Systems to Security Categories, available at <http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-60v1r1.pdf> (accessed 21st August, 2017).
 11. Volume II: Appendices to Guide for Mapping Information and Information Systems to Security Categories, available at <http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-60v2r1.pdf> (accessed 21st August, 2017).
 12. NIST Special Publication 800-53, Revision 4, 'Security and Privacy Controls for Federal Information Systems and Organizations', available at <http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf> (accessed 21st August, 2017).
 13. NIST Special Publication 800-18, Revision 1, 'Guide for Developing Security Plans for Federal Information Systems', available at <http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-18r1.pdf> (accessed 21st August, 2017).
 14. NIST Special Publication 800-30, Revision 1, 'Guide for Conducting Risk Assessments', available at <http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-30r1.pdf> (accessed 21st August, 2017).
 15. NIST Special Publication 800-30, Revision 1, 'Guide for Conducting Risk Assessments', available at <http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-30r1.pdf> (accessed 21st August, 2017).
 16. NIST Special Publication 800-18, Revision 1, 'Guide for Developing Security Plans for Federal Information Systems', available at <http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-18r1.pdf> (accessed 21st August, 2017).
 17. NIST Special Publication 800-37, Revision 1, 'Guide for Applying the Risk Management Framework to Federal Information Systems', available at <http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf> (accessed 21st August, 2017).